Introduction to Cloud Computing

..... for Enterprise Users

Lew Tucker, Ph.D.
CTO, Cloud Computing
Sun Microsystems, Inc.

Cloud computing means different things to different people

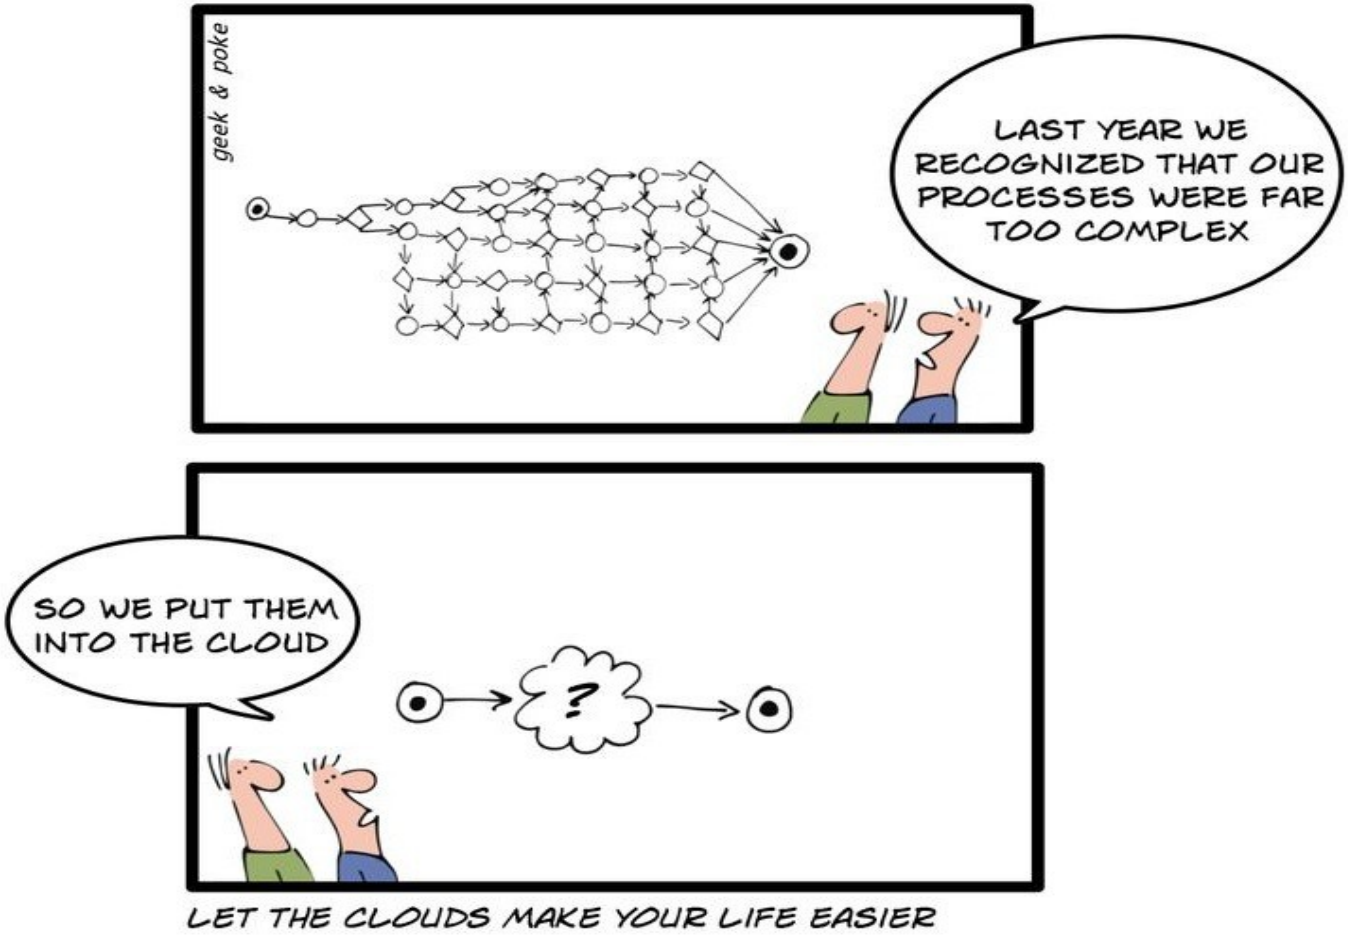

… and covers a lot of territory

Cloud Computing

is it the future?

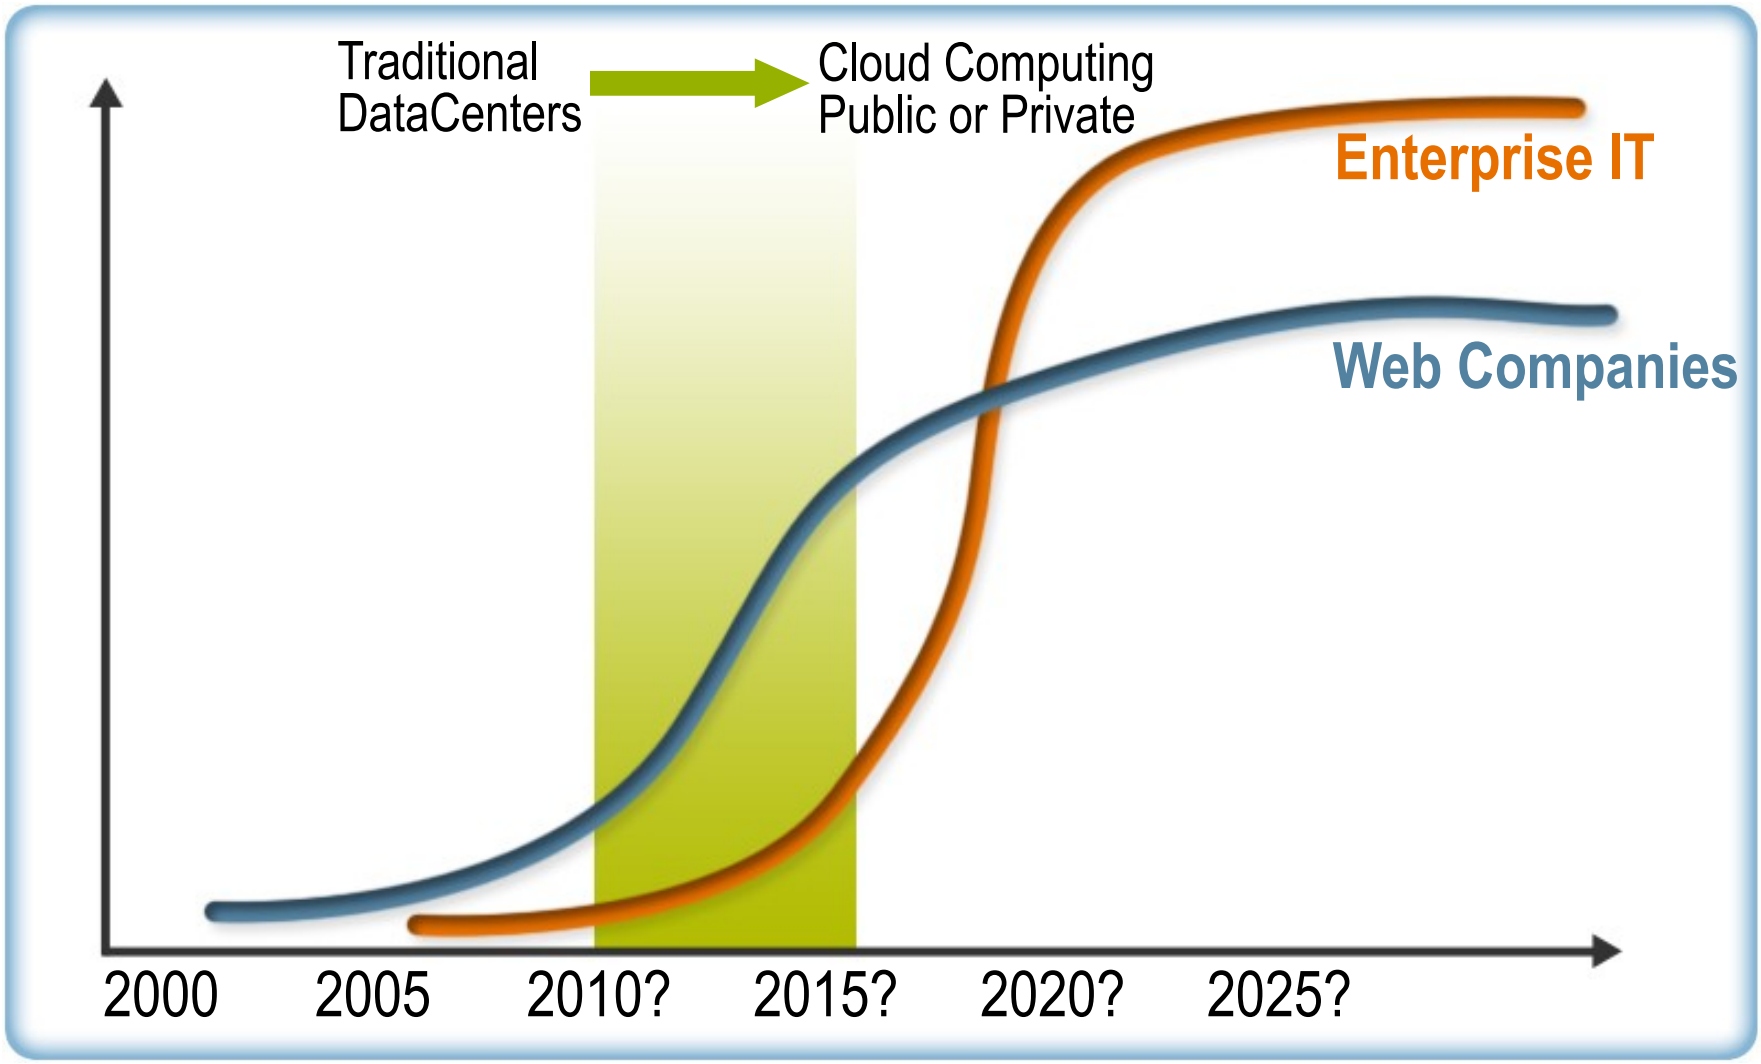

Alternative to traditional data centers

"Let me be very clear here:

I really don't want to operate datacenters anymore...

We'd rather spend our time giving our customers great service and writing great software rather than managing physical hardware,"

Don MacAskill, CEO, Smugmug

Definition by NIST

Cloud computing is a model for on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction.

http://csrc.nist.gov/groups/SNS/cloud-computing/index.html

Cloud Computing Attributes

Benefits: Efficiency, Flexibility and Speed

Economics

Pay As-You-Go

Op-ex vs. Cap-ex

Virtualization

Developer Centric

Rapid,
Self Provisioning

Faster Deployment

API-Driven

Flexibility

Highly Elastic

On Demand

Scalable Services

Number of players rapidly expanding

Cloud Computing 101

Cloud Computing Models

Software as a Service

Applications on-demand

Platform as a Service

Developer platform for creating applications

Infrastructure as a Service

Storage and compute capabilities offered as a service

Natural Evolution of the Web

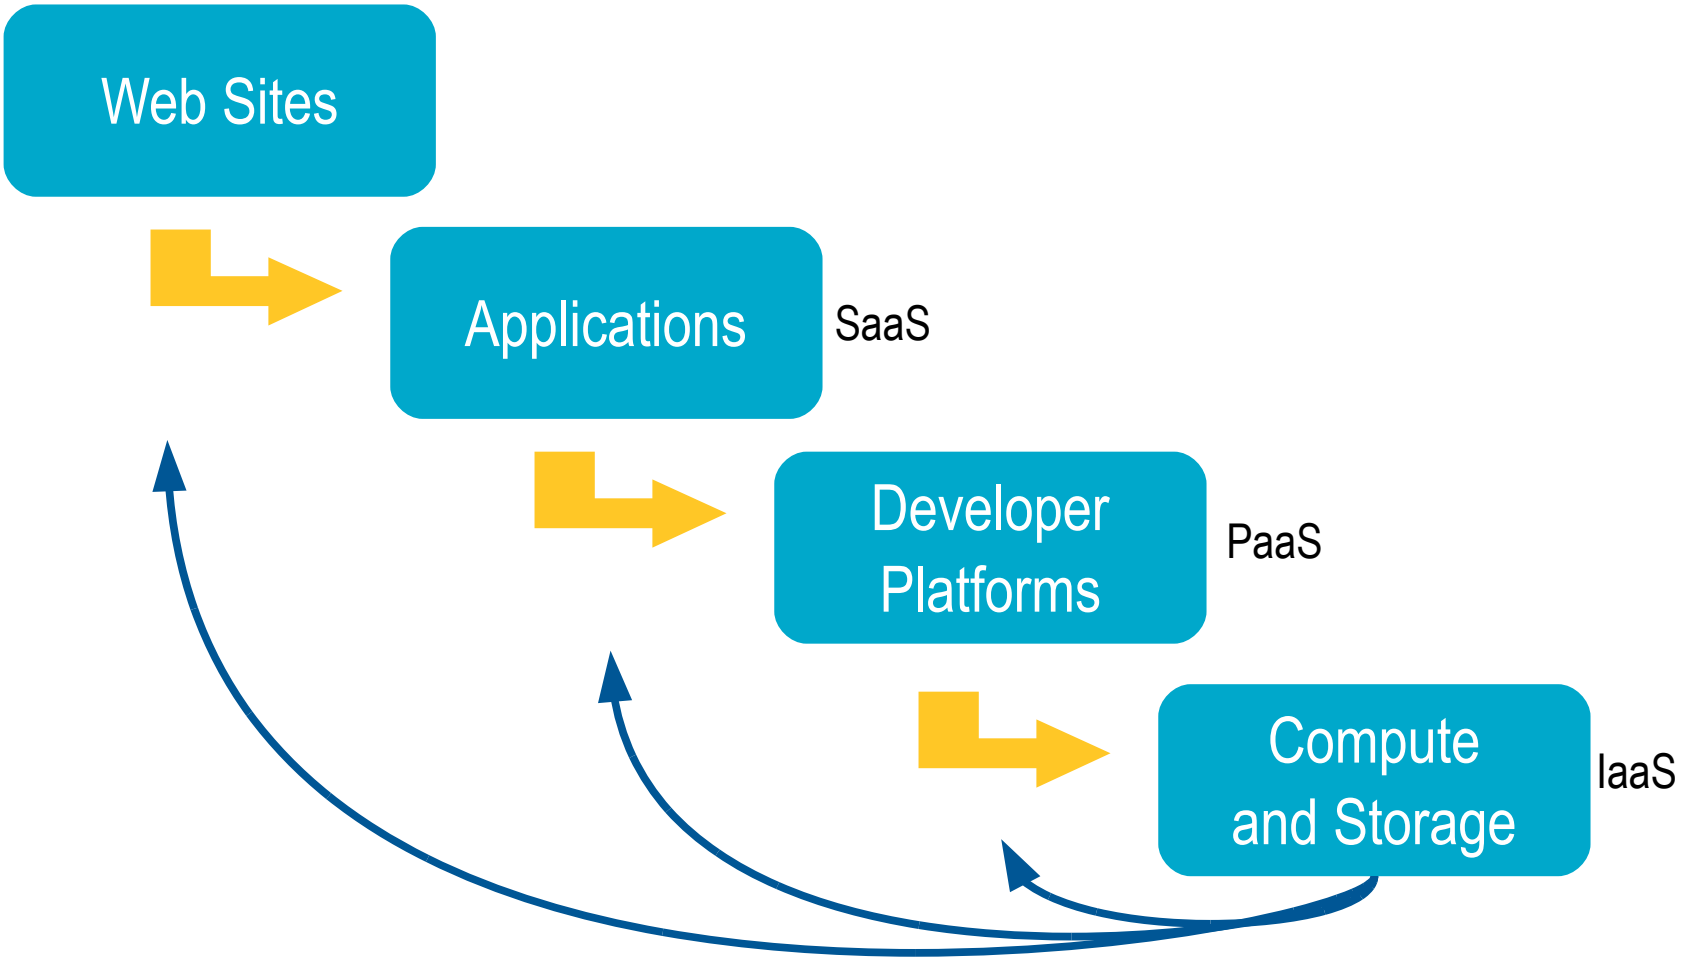

Software as a Service (SaaS)

Applications on demand:

- Subscription-based, multi-tenant, nothing to download or manage

- Google Apps (docs, email)
- Microsoft Exchange Online
- Yahoo Mail
- TurboTax Online
- Salesforce.com
- NetSuite
- Oracle CRM On-Demand
- Cisco WebEx Weboffice

Platform as a Service (PaaS)

On-demand develop and deploy apps

- Unique programming model, auto-scaling
- Often both a platform and a channel

- Google AppEngine
- Force.com (salesforce.com)
- Netsuite Business OS
- Heroku
- Aptana Cloud Connect
- Facebook

Infrastructure as a Service (IaaS)

On-demand virtual infrastructure

- Lowest level, most general, self-provisioning
- Unlimited managed resources

- Amazon AWS (EC2, S3, SQS)
- Microsoft Azure
- RackSpace Cloud
- Savis
- Terremark
- Joyent

Two more core concepts

- Virtual Machine Images
  - Complete, pre-configured, image of application and OS
  - Pre-packaged or built by user
- Cloud APIs
  - Programmic way to provision and manage compute, storage, and network resources
  - Access to scalable services (S3, SimpleDB)
  - Work underway to standardize for interoperability

Amazon's AWS

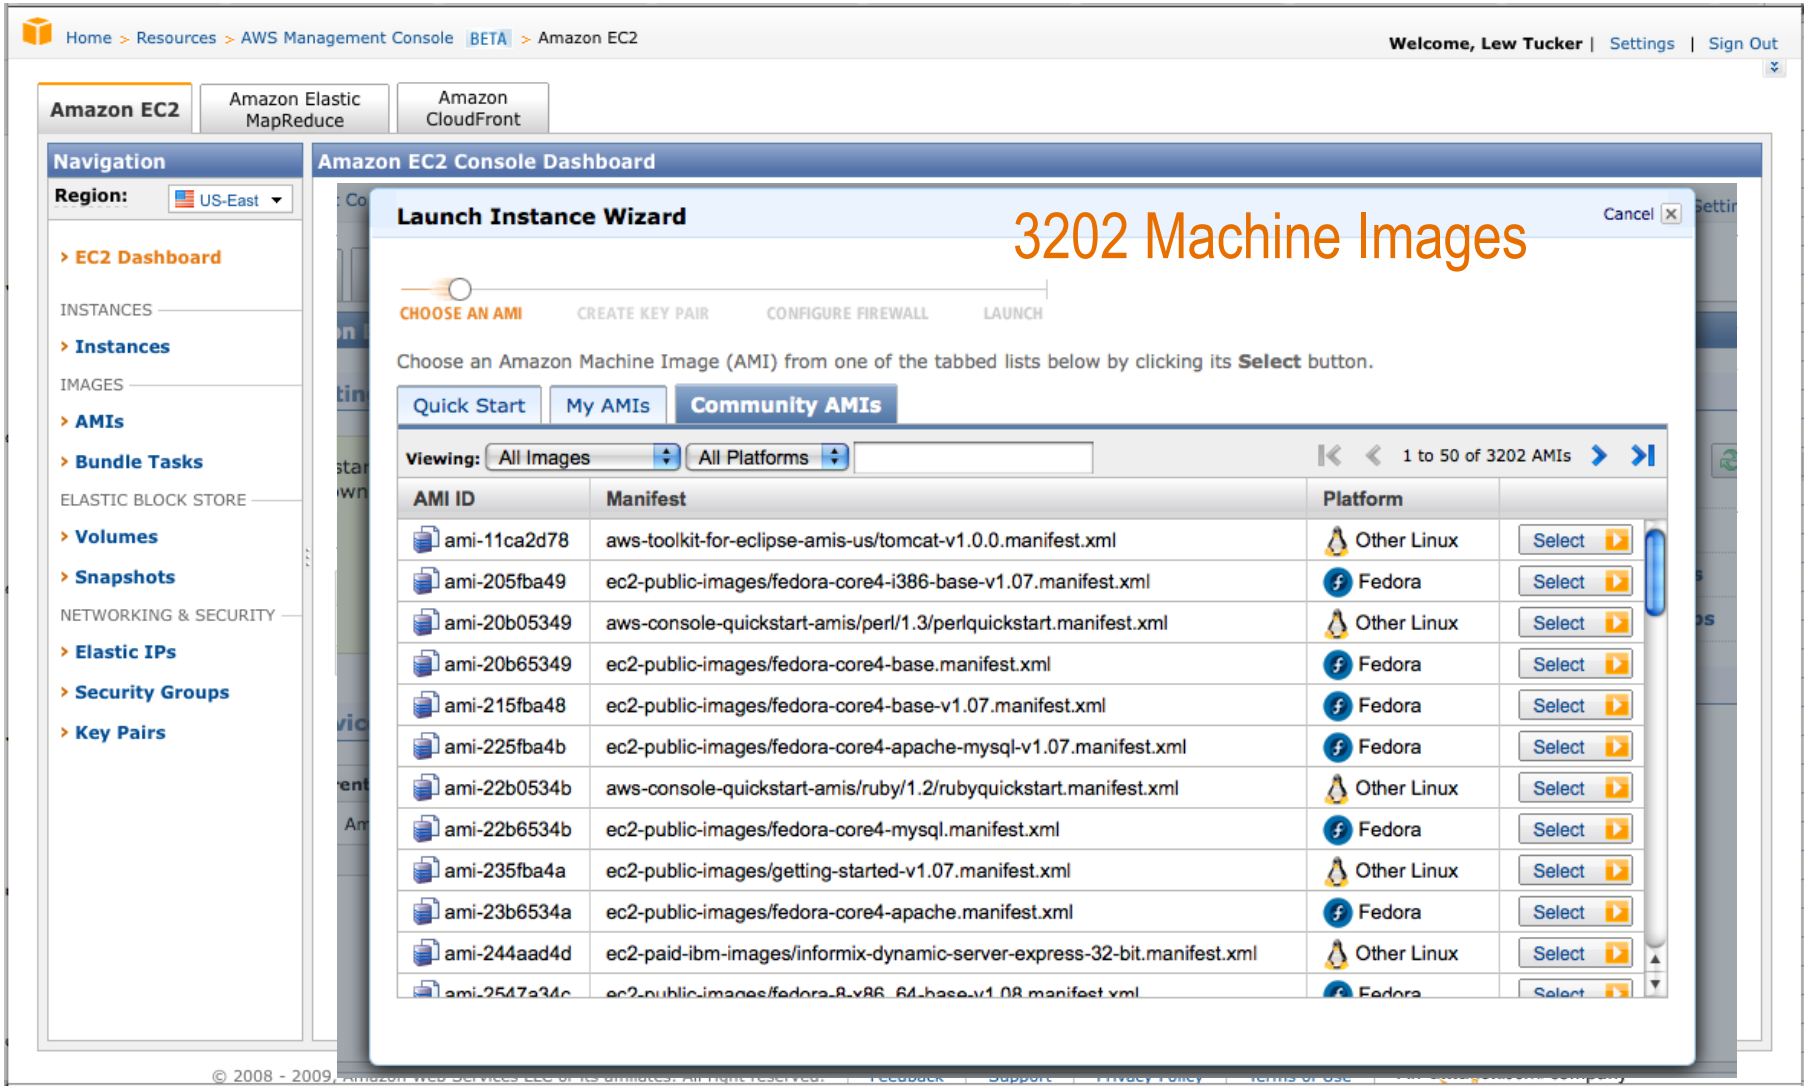
Home > Resources > AWS Management Console BETA > Amazon EC2
Welcome, Lew Tucker | Settings | Sign Out
Amazon EC2
Amazon Elastic MapReduce
Amazon CloudFront
Navigation
Region: US-East
EC2 Dashboard
INSTANCES
Instances
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
NETWORKING & SECURITY
Elastic IPs
Security Groups
Key Pairs
Amazon EC2 Console Dashboard
Launch Instance Wizard
Cancel
3202 Machine Images
CHOOSE AN AMI
CREATE KEY PAIR
CONFIGURE FIREWALL
LAUNCH
Choose an Amazon Machine Image (AMI) from one of the tabbed lists below by clicking its Select button.
Quick Start
My AMIs
Community AMIs
Viewing: All Images
All Platforms
1 to 50 of 3202 AMIs
AMI ID
Manifest
Platform
ami-11ca2d78 aws-toolkit-for-eclipse-amis-us/tomcat-v1.0.0.manifest.xml Other Linux Select
ami-205fba49 ec2-public-images/fedora-core4-i386-base-v1.07.manifest.xml Fedora Select
ami-20b05349 aws-console-quickstart-amis/perl/1.3/perlquickstart.manifest.xml Other Linux Select
ami-20b65349 ec2-public-images/fedora-core4-base.manifest.xml Fedora Select
ami-215fba48 ec2-public-images/fedora-core4-base-v1.07.manifest.xml Fedora Select
ami-225fba4b ec2-public-images/fedora-core4-apache-mysql-v1.07.manifest.xml Fedora Select
ami-22b0534b aws-console-quickstart-amis/ruby/1.2/rubyquickstart.manifest.xml Other Linux Select
ami-22b6534b ec2-public-images/fedora-core4-mysql.manifest.xml Fedora Select
ami-235fba4a ec2-public-images/getting-started-v1.07.manifest.xml Other Linux Select
ami-23b6534a ec2-public-images/fedora-core4-apache.manifest.xml Fedora Select
ami-244aad4d ec2-paid-ibm-images/informix-dynamic-server-express-32-bit.manifest.xml Other Linux Select
© 2008 - 2009

Issues – trusting a service provider

- Data governance and application security
  - Who has access
  - Trust in security of a shared, multi-tenant environment
- Legal
  - Who can see my data, where is it?
  - Third-party involvement in discovery
  - Regulatory compliance
- Business
  - Reliability
  - Lock-in

Public vs Private Clouds

Public

Pay as you go, multi-tenant applications and services

Access virtually unlimited resources

Private

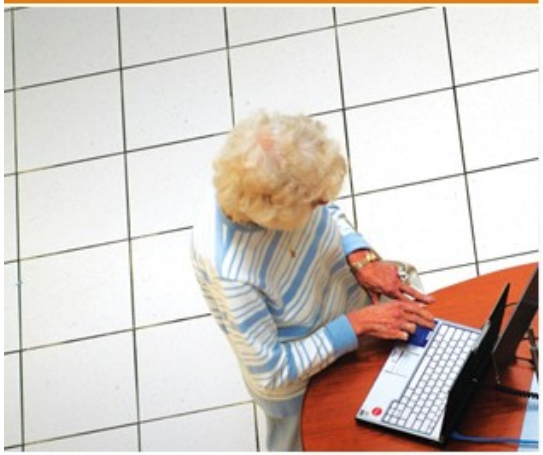

Cloud Computing model in a company's own datacenter

Resources directly owned but therefore constrained

Hybrid

Mixed usage of both public and private clouds, often integrated into the same application

Emerging usage of both public and private clouds in the enterprise

Public Cloud (service)

- Rapid provisioning of almost unlimited resources
- Pay only for what you need
- Opex vs Capex
- Departmental projects
- Analytics, Dev/test
- Customer-facing apps

Private Cloud (on-premise)

- Efficiencies of virtualization and data center automation
- Dynamic re-allocation of resources
- Reduction in operating costs
- Departmental self-service and chargeback

Software and services vendors for building private clouds

- VMware
- Citrix
- Eucalyptus
- Appistry
- Univa
- 3Tera
- Sun, Oracle, IBM, HP, Cisco
- Accenture, Deloite

IT-built cloud for internal customers using a shared pool of virtual resources

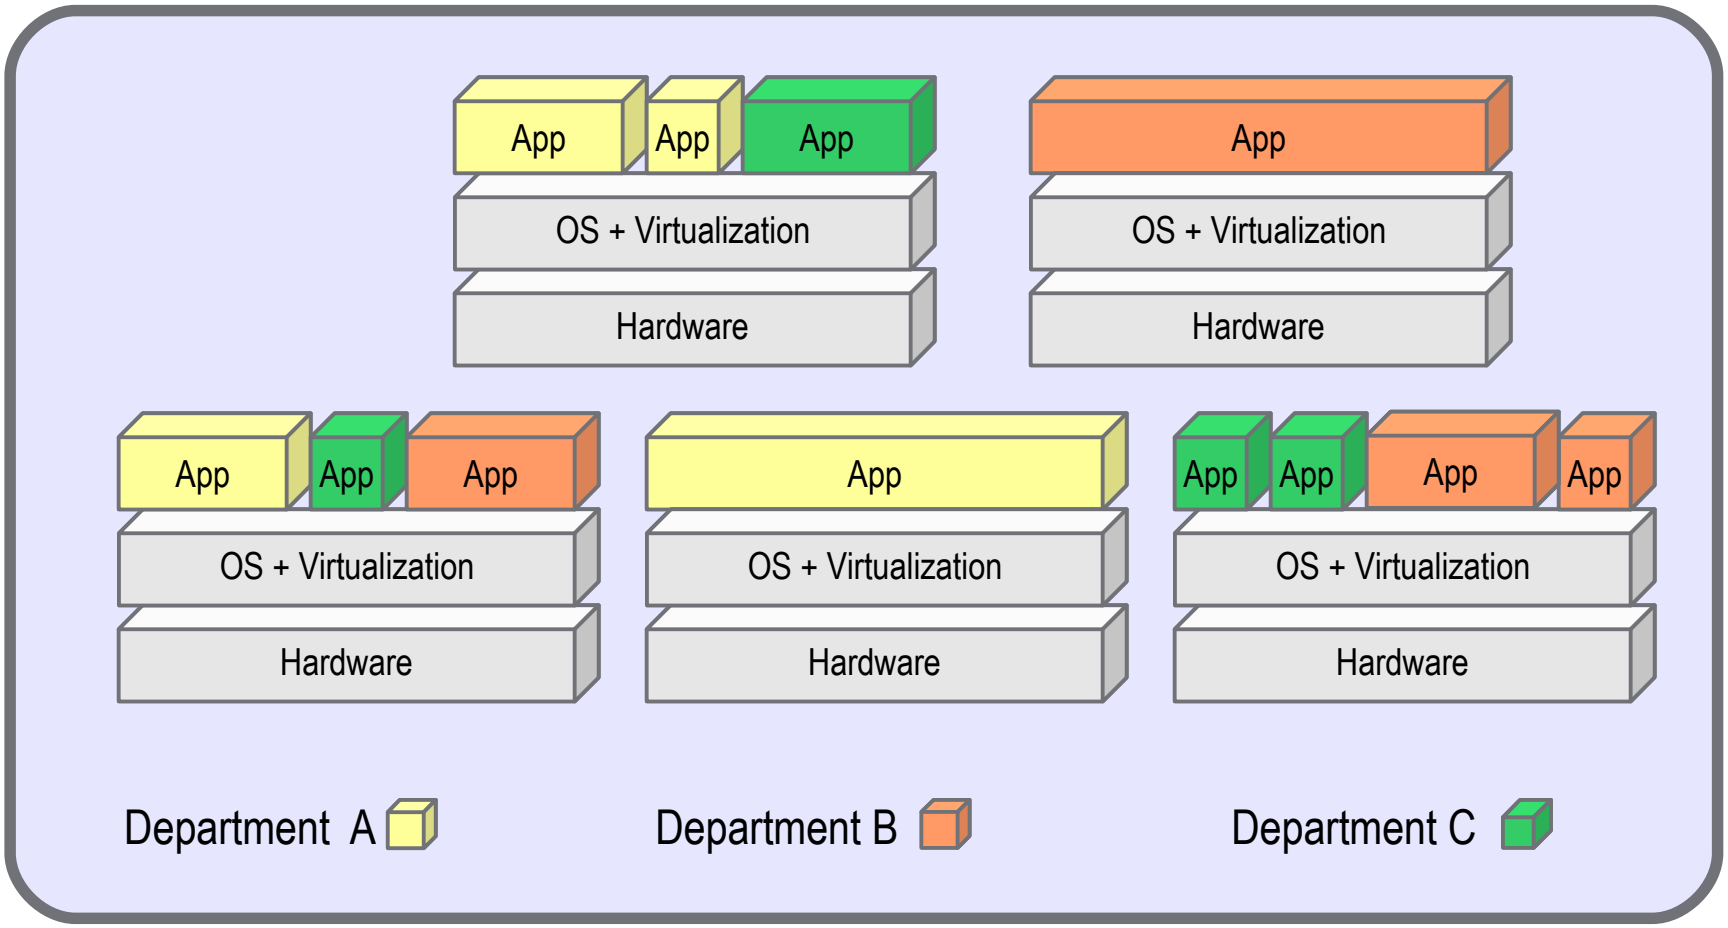

Expanded Role for IT

- Deliver the best technology solution for the business balancing cost, security, speed, user experience
  - Greater agility for business units through more options and self-service
- IT looks more like a business partner and service provider
  - Higher degree of automated system administration
  - Outsource many functions to pubic cloud providers
- Responsible for IT policy, security, best practices

Cloud Usage Patterns

Test and Development

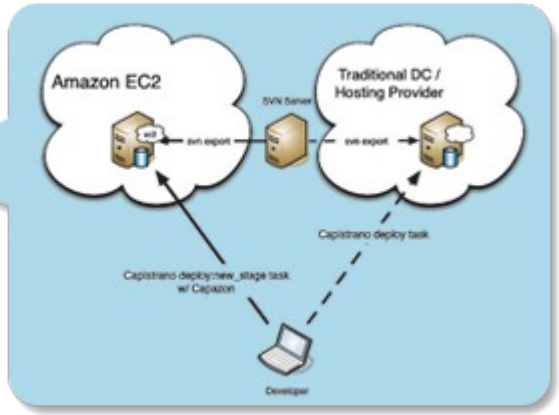

Functional Offload (Batch Processes – *TimesMachine*)

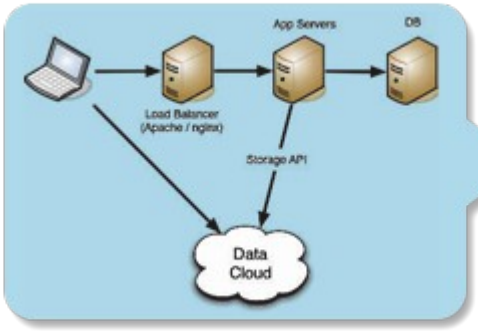

Functional Offload (Storage – *SmugMug*)

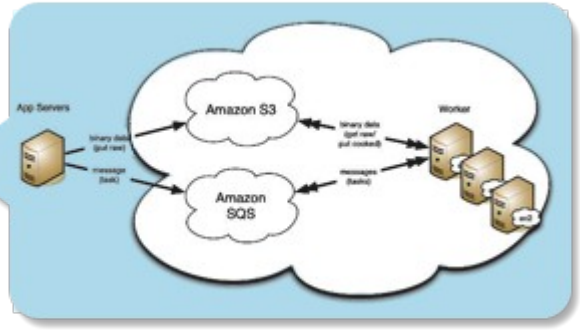

Cloud Bursting (Temporary Peak Load)

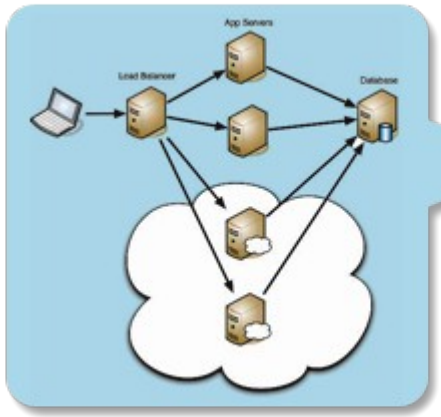

Web Service

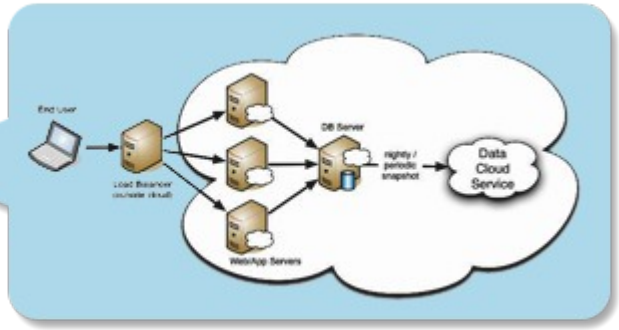

Welcome to Apps.gov

Apps.gov is your source for cloud computing applications designed to help your agency harness the power of today's technology. Whether it's Business or Productivity Applications, Cloud IT Services or Social Media solutions, Apps.gov is the place to get your government agency in the cloud.

What is Cloud Computing?

Want to learn more?

Watch this brief video for an overview of Cloud Computing to gain a better understanding of what it is and its benefits.

Watch the video now »

What type of solution do you need?

Business Apps

Your agency or service is complex and requires state-of-the-art software to get business done.

GSA Cloud Business Apps has a solution!

Cloud IT Services

Need a better solution to reduce cost and implement projects faster?

GSA Cloud IT Services has the answer!

Productivity Apps

You need to get things done and GSA is there to help you do just that.

GSA Cloud Productivity Apps has the tools!

Social Media Apps

Social media tools make it easier to discuss the things we care about and help us get the job done.

GSA Social Media Apps can help you get the word out!

GSA Cloud Computing Storefront

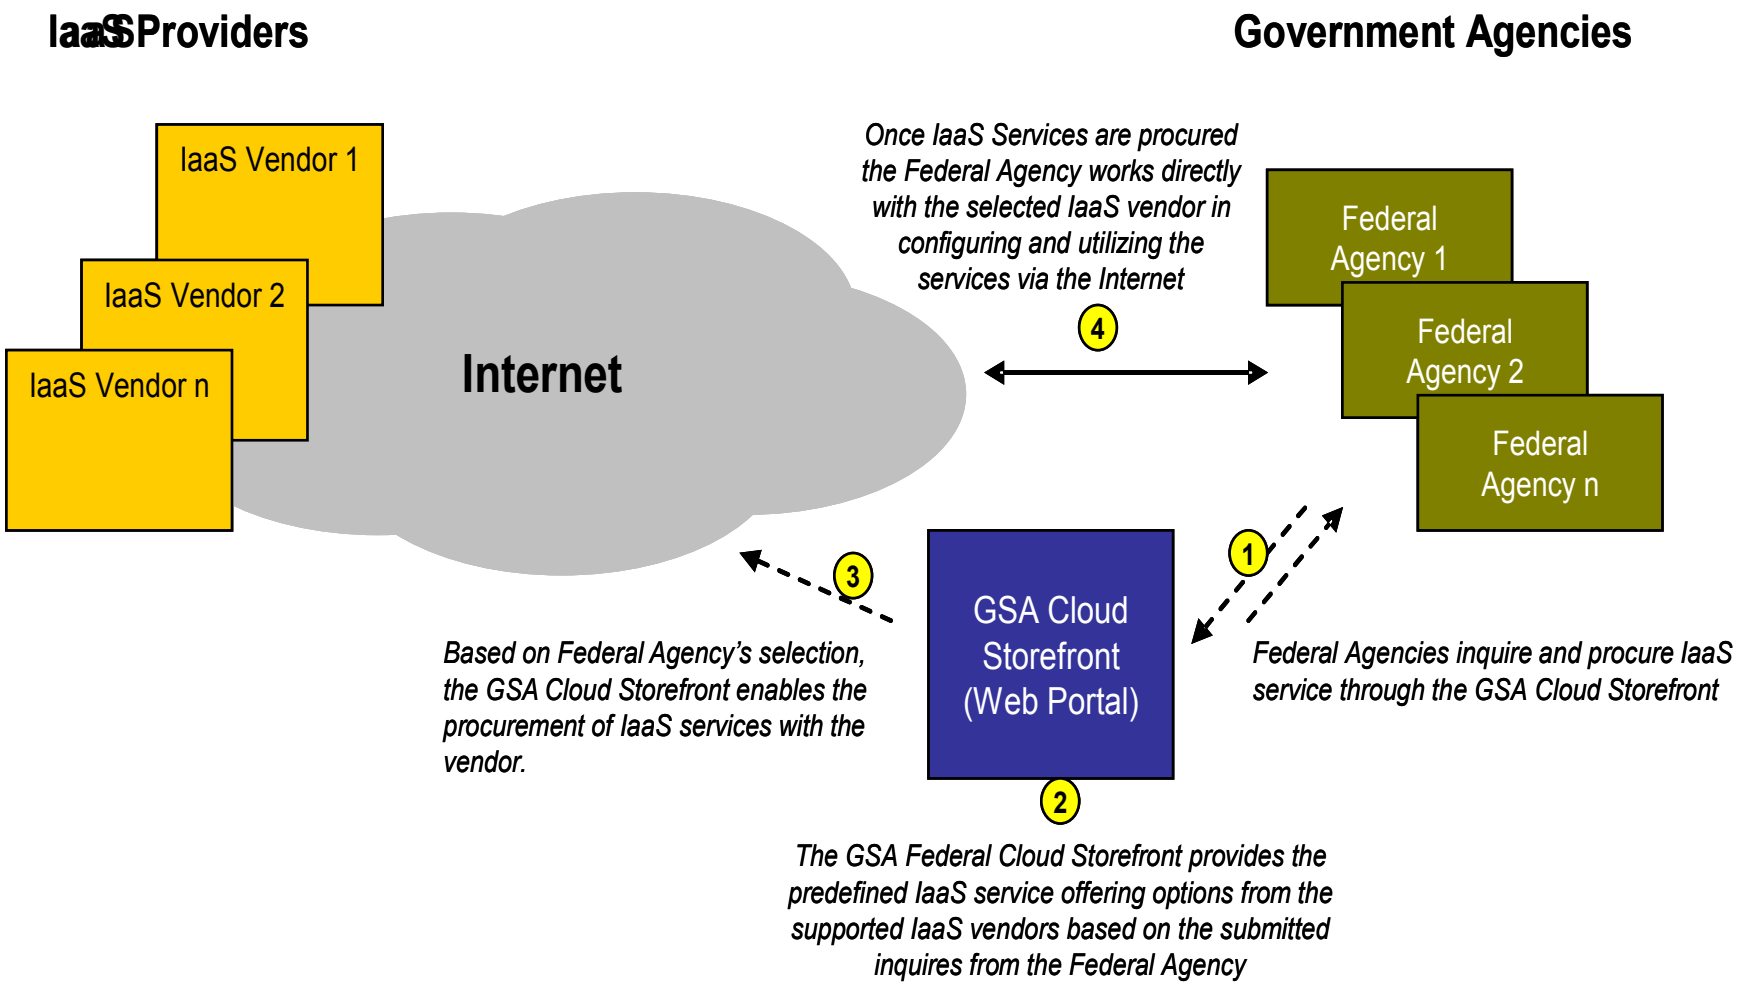

USA.gov and Data.gov

- Federal government's most high-profile websites
- Reduction in annual expenses $2.5M to $800k
- Deploy new apps in 24 hours
- Built on Terremark's Enterprise Cloud platform

USA.gov Government Made Easy

terremark NAP of the Americas

MedCommons: health records sharing

- Health records services provider allowing consumers to share info using Amazon's AWS
- Leverages Amazon's billing systems
- Subscription-based
- HIPPA compliant

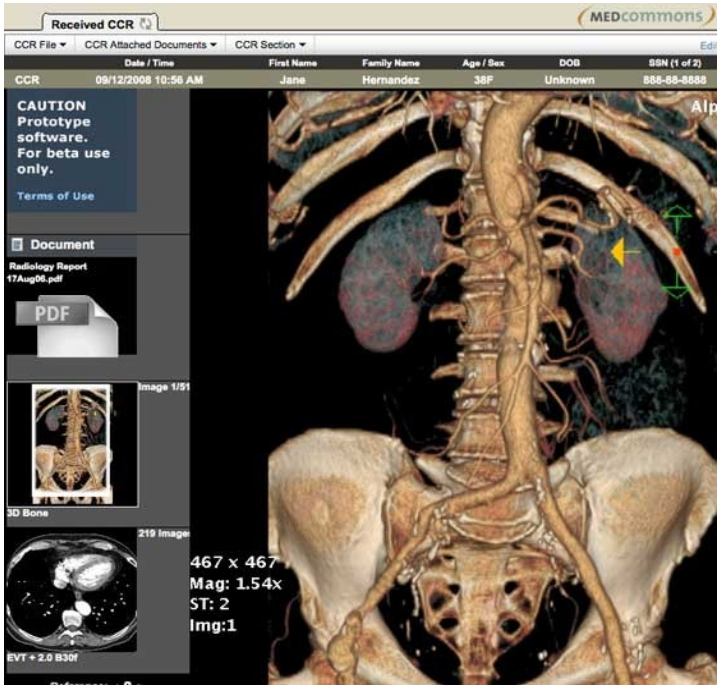

http://aws.amazon.com/solutions/case-studies/medcommons/

Washington Mutual: internal cloud

- Goal: reduce waste, complexity, and costs
- Multi-phase approach
  - Started with existing physical servers (< 10% utilization)
  - Moved to a fixed number of virtual servers (VMware)
  - Dynamic allocation of VMs (30-day increments)
  - Self-service, compute on demand, highly automated ops
- Experience
  - Utility computing 40-70% more cost effective
  - Unit costs of individual VMs down 60% in 18 months
  - Deployment now < 5 days

Washington Mutual - advice

- Standardize and simplify all offerings
- Use server consolidation to drive critical mass
- Automate and deploy tools to drive transparency for platform users

"I see the difference between utility and cloud-based computing is that the enterprise cloud really focuses on true, on-demand compute. When somebody needs it, they get it. And if they need it for one day, they get it for just one day" - *Barton Warner, VP*

Internal clouds as the next generation of data center management

- Builds on virtualization and server consolidation
- Improves cost-effectiveness of IT workloads
- Highly automated, removes operators from provisioning
- Provides departmental self-service
- Improves business agility up to the constraint of fixed resources (is this still a cloud?)
- Potentially capable of “bursting” out to public clouds

New Data and Management Economics

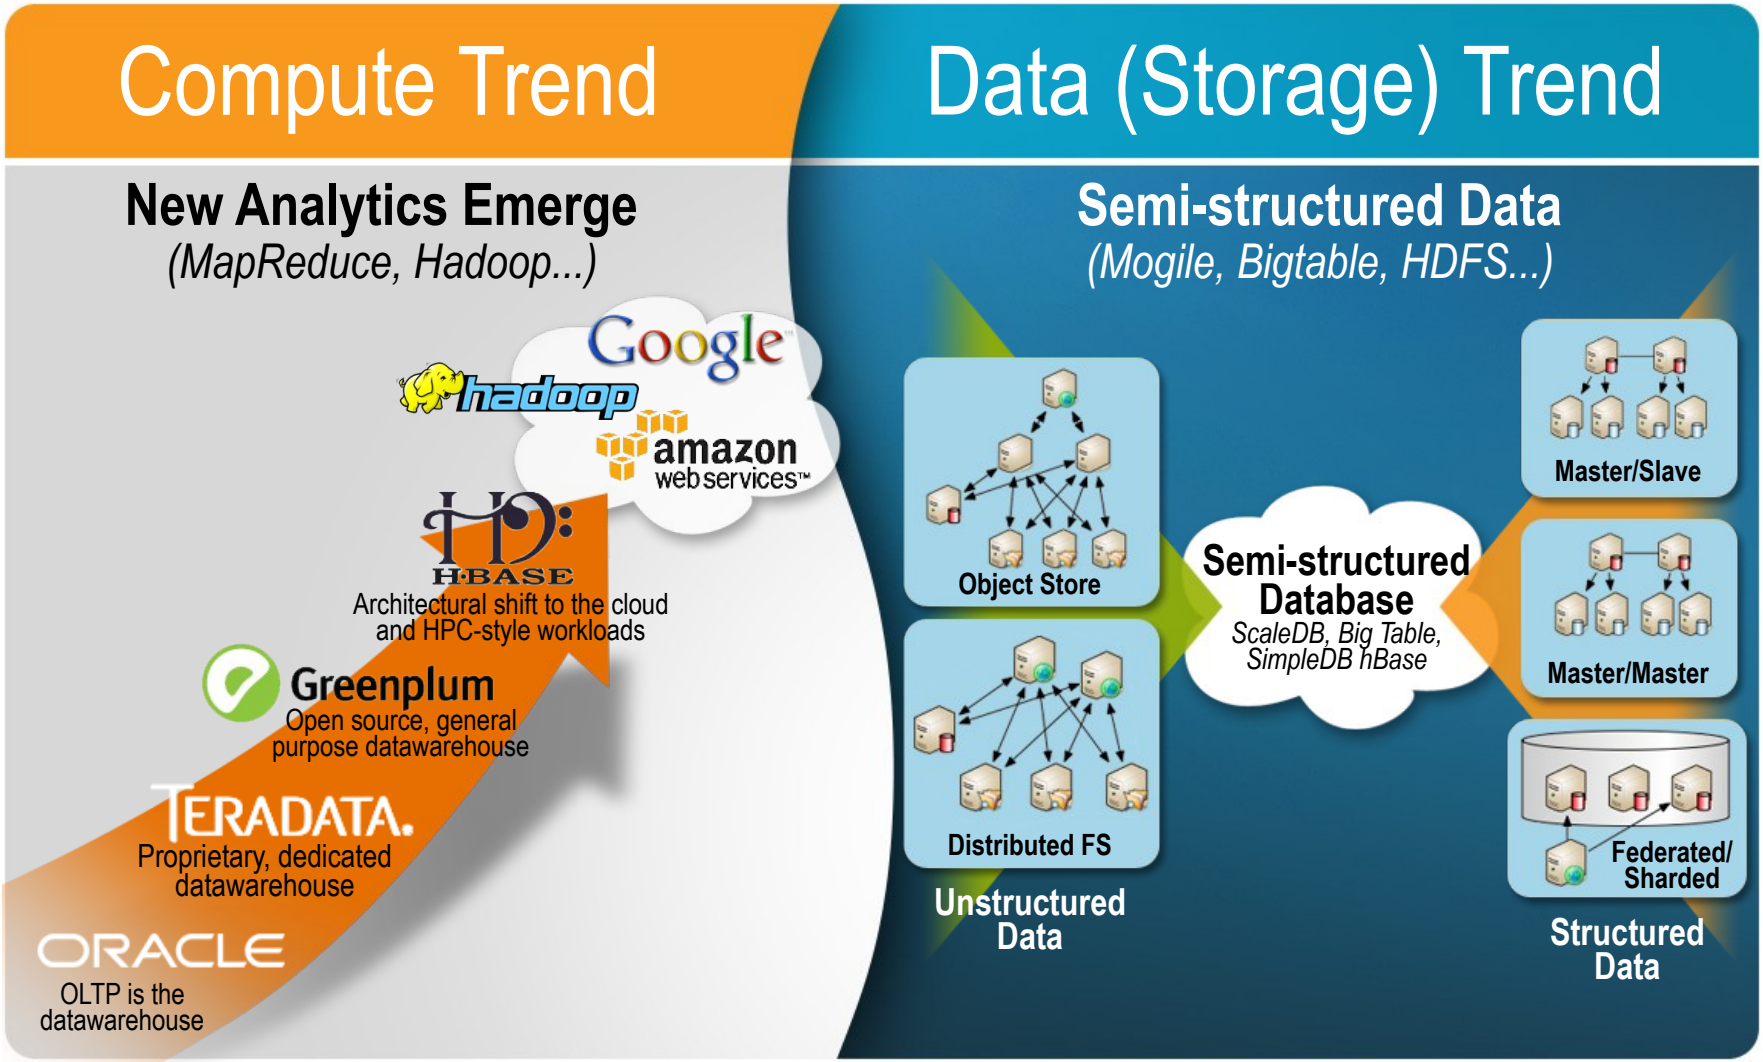

GeoEye: satellite and aerial imagery

- Satellite and aerial imagery and geospatial products for national security and commercial customers
- Combine on-premise Appistry cloud with Amazon's S3 service
- Appistry CloudIQ uses 50+ servers for image processing
- Amazon's S3 serves up imagery to customers

GeoEye

Appistry

amazon web services™

http://www.appistry.com/customers/profile/geoeye

Agile Analytics in the Cloud

Enterprise software vendors team up to provide a business intelligence solution that any business can setup in a matter of minutes.

It's Not Just About Cheap Computing

Best way is to simply try it yourself

Cloud Computing Resources

- Sun Resources
  - https://www.sun.com/offers/details/cloud_computing_primer.xml
  - https://www.sun.com/offers/details/CloudComputing.xml
  - http://www.sun.com/solutions/cloudcomputing/perspectives.jsp
  - http://kenai.com/projects/suncloudapis
  - http://developers.sun.com/cloud/
  - http://blogs.sun.com/ec2/entry/hardened_opensolaris_2008_11_on
  - http://kenai.com/projects/s3-crypto/pages/Home
  - http://kenai.com/projects/zfs-backup-to-s3/pages/Home
  - http://www.sun.com/service/cloud/
  - http://wikis.sun.com/display/cloud/Patterns
  - http://wikis.sun.com/display/VeriScale/Home
- Other Resources
  - http://groups.google.com/group/cloud-computing?lnk=
  - http://groups.google.com/group/cloudforum?lnk=
  - http://searchcloudcomputing.techtarget.com/
  - http://cloudcomputing.sys-con.com/
  - http://cloudbook.net/
  - http://www.eucalyptus.com/